

Distributed Systems

2013 Exam 3 Review

Paul Krzyzanowski

Rutgers University

Fall 2013

Fall 2013 - Question 1

How does a clustered file system differ from a distributed file system (e.g., NFS, SMB, AFS, Coda)?

- **Clustered file system**
 - Block-level access to storage. File system implemented at the client OS.
- **Distributed file system**
 - Remote access to files.

Fall 2013 - Question 2

Why is it important to use consistent hashing in a distributed hash table?

Note: The question is not asking you to define consistent hashing.

- To avoid moving an excessive amount of data among nodes.
- With consistent hashing, only some data from a neighboring node(s) has to be moved.

Fall 2013 - Question 3

Explain the difference between a public key and symmetric algorithm.

- **Symmetric encryption**: same key is used for encryption and decryption.
- **Public key encryption**: a pair of related keys, K_1 and K_2 , is used for encryption and decryption.
 - If K_1 is used to encrypt, then K_2 must be used to decrypt
 - If K_2 is used to encrypt then K_1 must be used to decrypt.

Fall 2013 - Question 4

Alice sends Bob her X.509 digital certificate.

Bob validates the certificate successfully.

How does he now validate that he is indeed communicating with Alice?

- By possessing Alice's certificate, Bob has her public key.
- He needs to prove that Alice has the corresponding private key.
 1. Bob generates a random string (nonce) and sends it to Alice.
 2. Alice encrypts it with her private key and sends the result to Bob.
 3. Bob decrypts the received message using Alice's public key (from her certificate). If the result matches the nonce, he is convinced.

Fall 2013 - Question 5

A digital signature or message authentication code can protect us from certain:

- a) Byzantine faults.
- b) Fail stop faults.
- c) Fail silent fault.
- d) Fail restart faults.

- Can allow us to detect if a message is modified
- But ... does not detect retransmission

Fall 2013 - Question 6

Chubby's fault tolerance model is:

- a) Active-active.
- b) **Active-passive.**
- c) Triple modular redundancy (TMR).
- d) Five-way modular redundancy (5-MR).

- Active-passive = one server processes requests and propagates state to replicas

Fall 2013 - Question 7

For a system to be k -fault tolerant in the presence of faults that may be either byzantine or fail-silent, you need this many components:

- a) $k + 1$
- b) $2(k+1)$
- c) $2k + 1$
- d) $k^2 + 1$

- k components may produce faulty results
- $k+1$ good ones will force a majority vote
- Total components = $k + (k + 1) = 2k + 1$

Fall 2013 - Question 8

An asynchronous network makes it difficult to design a system that will:

- a) Determine that a computer is not communicating.
- b) Determine the ordering of events.
- c) Identify the origin of a message.
- d) Distinguish causal messages from concurrent messages.

- No upper bound on message transit
- Unsure of whether a message is delayed (or lost) or has not been sent

Fall 2013 - Question 9

Quorum in a cluster is important to ensure that:

- a) More than one group of computers do not create their own cluster.
- b) There is sufficient computing power available for the task.
- c) All computers in the cluster are alive.
- d) All computers in the cluster have a backup.

- Prevent split brain

Fall 2013 - Question 10

A heartbeat is used to:

- a) Detect dead computers in a cluster.
- b) Synchronize operations in a cluster.
- c) Provide high-speed communication links within a cluster.
- d) Propagate configuration changes throughout the cluster.

Fall 2013 - Question 11

A load balancer is least useful for:

- a) Migrating processes from one computer to another.
- b) Distributing requests among a pool of servers.
- c) Fault tolerance.
- d) Allowing an administrator to take a server out of a cluster for upgrades with no downtime.

- A load balancer distributes requests
- It does not support the migration of workloads

Fall 2013 - Question 12

Map workers and Reduce workers in a Google MapReduce cluster use this failover model:

- a) Cold
- b) Warm
- c) Hot
- d) Passive

- Process restarts or takes over with:
 - Cold = no saved state of the computation
 - Warm = state from the last checkpoint (e.g., Pregel)
 - Hot = no lost state (e.g., Chubby)

Fall 2013 - Question 13

A Google cluster comprises computers that are selected for the:

- a) **Best energy efficiency to performance ratio.**
- b) Maximum CPU performance.
- c) Fastest local storage.
- d) Smallest size.

Fall 2013 - Question 14

Looking up the address and port of a server at the start of a client process is an example of:

- a) Static binding.
- b) **Early binding.**
- c) Late binding.
- d) Delayed binding.

- Static binding = hard-coded binding
- Early binding = *a priori* lookup
- Late binding = resolve immediately before use
- Delayed binding = ???

Fall 2013 - Question 15

The Domain Name System (DNS) is built with a distributed lookup that uses:

- a) A central coordinator.
- b) Flooding.
- c) Referrals.
- d) A distributed hash table.

Fall 2013 - Question 16

An overlay network is a:

- a) Set of connections that define a spanning tree to ensure there are no cycles.
- b) Private network of high-speed connections that overlays part of the public Internet.
- c) Wireless network that overlays the wired Internet.
- d) Graph whose edges identify nodes that know about each other.

Fall 2013 - Question 17

CAN, the Content-Addressable Network is a peer-to-peer storage system that:

- a) Allows a client to locate an object by any of its content instead of a key.
- b) Enables a client to locate an object via multiple keys, one per axis in each dimension.
- c) Transforms a key into an address of the server holding the corresponding object.
- d) Hashes a single key into multiple axes, one per dimension.

- (a) No. We look up a key
- (b) No. Just one key
- (c) Each host holds keys that hash into a range of values but you cannot transform the key into an address of a host
- (d) A key is hashed once per dimension to identify its place in the grid

Fall 2013 - Question 18

A finger table in a Chord node is:

- a) A table of frequently used *key*→*node* mappings.
- b) A tree structure that enables a node to find any other node in $O(\log N)$ table reads.
- c) A table with each element, i , representing a node that is i hops away.
- d) A table with each element, i , representing a node that is 2^i hops away.

Fall 2013 - Question 19

Dynamo's structure is most similar to:

- a) Bigtable.
- b) Flooding.
- c) Chord.
- d) CAN.

- Logical ring of nodes.
- Each virtual node holds a contiguous range of hash values

Fall 2013 - Question 20

Unlike Bigtable, with Amazon Dynamo:

- a) Keys are sorted alphabetically to support iteration.
- b) An object is identified by exactly one key.
- c) **Two processes may write conflicting updates.**
- d) All requests pass through a coordinator.

(a) This is a property of Bigtable

(b) Both Dynamo and Bigtable use a single key

(d) Neither Dynamo nor Bigtable send requests through a coordinator

- Multiple processes may end up writing conflicting values to the same key with Dynamo. Vector timestamps identify concurrent updates.

Fall 2013 - Question 21

Virtual nodes in Amazon Dynamo are designed to:

- a) Improve fault tolerance due to the replication of nodes.
- b) Increase the requests the system can handle since many virtual nodes can be managed by one physical node.
- c) **Improve load distribution when adding or removing nodes.**
- d) Create an overlay network that arranges nodes into a logical ring.

(a) Virtual nodes are not replicated. Data is replicated among physical nodes. Virtual nodes help with balancing load if a node dies.

(b) The performance is a function of the capacity of physical nodes.

(d) The logical ring is there with or without the use of virtual nodes.

- *A newly available node accepts a roughly equivalent amount of load from each of the other available nodes.*

Fall 2013 - Question 22

Akamai uses DNS to resolve a domain name to:

- a) **The nearest server that has the desired cached content.**
- b) A load balancer that then forwards the request to any available caching server.
- c) A coordinator that will analyze the request and forward it to the nearest caching server.
- d) The original server, which then sends an HTTP REDIRECT message to the nearest caching server.

- Goal is to find the best (nearest/fastest) server with content
- (b) No – a request is not forwarded to *any* caching server
- (c) No – there is no coordinator that analyzes requests
- (d) No – the original server is contacted only by caching servers if no caching server has the content

Fall 2013 - Question 23

A hash function is useful in the generation of a:

- a) Nonce.
- b) Symmetric key.
- c) Digital signature.
- d) Session key.

(a) Could be used but pointless: this is a random bunch of bits.

(b) Could be used but pointless: this is a random bunch of bits.

(c) Yes.

(d) Could be used but pointless: this is a random bunch of bits.

Fall 2013 - Question 24

SSL is an example of a:

- a) Symmetric key cryptosystem.
- b) Public key cryptosystem.
- c) Hybrid cryptosystem.
- d) Restricted cipher.

- Hybrid cryptosystem = public key for authentication and/or key exchange, symmetric for communication

Fall 2013 - Question 25

The Challenge Handshake Authentication Protocol (CHAP) tests to see if you know a:

- a) **Secret value.**
- b) Public key.
- c) Private key.
- d) Session key.

- Test knowledge of shared secret
(may be a secret key, PIN, or other data)

Fall 2013 - Question 26

Alice wants to talk to Bob and gets a ticket from a Kerberos server. The ticket is:

- a) Encrypted so only Alice can decode it.
- b) Encrypted so only Bob can decode it.
- c) Encrypted so only Alice and Bob can decode it.
- d) Not encrypted but contains a digital signature so that Alice and Bob can validate it.

Alice gets two things from Kerberos:

1. A message encrypted for her containing a session key & Bob's contact
2. A ticket (sealed envelope) that is encrypted for Bob that contains the same session key and Alice's contact

Fall 2013 - Question 27

A random number generator is NOT useful for generating a:

- a) Nonce.
- b) Symmetric key.
- c) Digital signature.
- d) Session key.

(a) A nonce is a random bunch of bits.

(b) symmetric key is a random bunch of bits.

(d) A session key is used as a symmetric key and is a random bunch of bits.

- A digital signature is an encrypted hash of a message.

Fall 2013 - Question 28

With OpenID, a web site:

- a) Identifies the user but allows another site to authenticate the user.
- b) Sends the user's ID and password to another site to validate them.
- c) Requests the user's password from the Identity Provider and uses that to authenticate the user.
- d) Provides anonymous login capabilities since the site never knows the identity of the authorized user.

- (a) OpenID delegates authentication to another party (Identity Provider)
- (b) No. The password is entered with the Identity Provider
- (c) No. The password is never sent back
- (d) No. The user is identified (might be non-binding)

Fall 2013 - Question 29

OAuth differs from OpenID in that it:

- a) Handles authentication while OpenID is responsible for identification.
- b) Sends the user to another site to approve access to specific services at that site.
- c) Uses the HTTP REDIRECT to send the user to a third-party site.
- d) Provides a centralized server to manage information access requests for all users.

(a) OpenID is responsible for authentication. OAuth is responsible for service authorization

(b) YES

(c) Yes, but so does OpenID

(d) No. Both OpenID and OAuth are decentralized

The End